arXiv:1511.07826v2 [cs.DS] 1 Dec 2015

Lift-and-Round to Improve Weighted Completion Time on Unrelated Machines

Nikhil Bansal* Aravind Srinivasan† Ola Svensson‡

December 2, 2015

Abstract

We consider the problem of scheduling jobs on unrelated machines so as to minimize the sum of weighted completion times. Our main result is a $(3/2 - c)$-approximation algorithm for some fixed $c > 0$, improving upon the long-standing bound of 3/2 (independently due to Skutella, *Journal of the ACM*, 2001, and Sethuraman & Squillante, *SODA*, 1999). To do this, we first introduce a new lift-and-project based SDP relaxation for the problem. This is necessary as the previous convex programming relaxations have an integrality gap of $3/2$. Second, we give a new general bipartite-rounding procedure that produces an assignment with certain strong negative correlation properties.

Keywords: Approximation algorithms, semidefinite programming, scheduling

1 Introduction

We consider the classic problem of scheduling jobs on unrelated machines to minimize the sum of weighted completion times. Formally, a problem instance consists of a set $J = \{1, 2, \ldots, n\}$ of n jobs and a set M of m machines; each job $j \in J$ has a weight $w_j \geq 0$ and it requires a processing time of $p_{ij} \geq 0$ if assigned to machine $i \in M$. The goal is to find a schedule that minimizes the weighted completion time, that is $\sum_{j \in J} w_j C_j$, where C_j denotes the completion time of job j in the schedule constructed.

Total completion time and related metrics such as makespan and flow time, are some of the most relevant and well-studied measures of quality of service in scheduling and resource allocation problems. While total completion time has been studied since the 50's [32], a systematic study of its approximability was started in the late 90's by [25]. This led to a lot of activity and progress on the problem in various scheduling models and settings (such as with or without release dates, preemptions, precedences, online arrivals etc.). In particular, we have now a complete understanding of the approximability in *simpler* machine models, such as identical and related machines. For these settings, non-trivial approximation schemes were developed more than a decade ago, e.g., in [1, 31, 9]. The more general unrelated machine model behaves very differently and is significantly more challenging. Perhaps because of this, its study has led to the development of many new techniques, such as interesting LP and convex programming formulations and rounding techniques [25, 11, 16, 26, 15, 30] (see also the survey by Chekuri and Khanna [10]).

*Department of Mathematics and Computer Science, TU Eindhoven, Netherlands. Email: n.bansal@tue.nl. Supported by NWO Vidi grant 639.022.211 and ERC consolidator grant 617951.

†Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, USA. Email: srin@cs.umd.edu. Supported in part by NSF Awards CNS-1010789 and CCF-1422569, and a research award from Adobe, Inc.

‡School of Computer and Communication Sciences, EPFL, Switzerland. Email: ola.svensson@epfl.ch. Supported by ERC Starting Grant 335288-OptApprox.

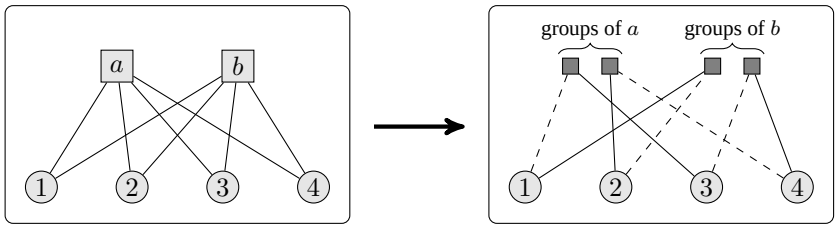

Figure 1: A simple example motivating the novel rounding algorithm with strong negative correlation.

In spite of these impressive developments, it remains a notorious problem to understand the approximability of total weighted completion time in the unrelated machines setting. On the positive side, Schulz and Skutella [26] gave a $(3/2 + \epsilon)$-approximation based on a time-indexed LP formulation, improving upon the previous works of [24, 17]. Another $3/2$-approximation was obtained by Skutella [30] and independently by Sethuraman and Squillante [28] based on a novel convex-programming relaxation. On the other hand, Hoogeveen et al. [18] showed that the problem is APX-hard, but the hardness factor was very close to 1. The natural question of whether a $(3/2 - c)$-approximation exists for the problem for some $c > 0$ has been proposed widely [10, 26, 20, 34, 27]. In particular, it appears as *Open Problem 8* in the well-known list [27] due to Schuurman and Woeginger of the "top ten" problems in scheduling. Moreover, Srividenko and Wiese conjectured [34] that the configuration LP for this problem has an integrality gap that is strictly less than $3/2$. The unrelated machines setting is one of the most general and versatile scheduling models that incorporates the heterogeneity of jobs and machines. But besides the practical motivation, an important reason for interest in the problem is that historically, the exploration of various scheduling problems in the unrelated machines model has been a rich source of several new algorithmic techniques [22, 33, 7, 8, 5, 4, 13, 23, 21, 6].

Our Results: Our main result is such an improved algorithm. In particular, we show the following.

Theorem 1.1. *There is a $(3/2 - c)$-approximation algorithm for minimizing the total weighted completion time on unrelated machines, for some $c \geq 10^{-7}$.*

Remark: We do not try to optimize the constant c too much, preferring instead to keep the exposition as simple as possible. However, it does not seem likely that our analysis would yield $c < 10^{-2}$.

The result is based on two key ideas: (i) a novel SDP relaxation for the problem obtained by applying one round of lift-and-project to the standard LP formulation and (ii) a new rounding algorithm to assign jobs to machines that reduces the "correlation" between the various jobs assigned to a machine.

Stronger Formulation: The stronger formulation is necessary: we show that the convex programming relaxation considered by [30, 28] has an integrality gap of $3/2$. Such families of instances do not seem to have been previously known [29]; we describe them in Section 2. In contrast, our stronger relaxation can be used to derive new and tighter lower bounds on the optimum value, by exploiting the PSD constraint on the underlying moment matrix.

We remark that the lower bounds we use to prove Theorem 1.1 can also be obtained using the configuration LP proposed by [34], which confirms their conjecture that the integrality gap of the configuration LP is also upper bounded by $(3/2 - c)$. However, we find the SDP formulation more natural as it is explicitly reveals the correlation information that we use.

The Rounding Algorithm: The solution to the SDP gives a fractional assignment of jobs to machines, which we need to convert to an integral assignment. Interestingly, all the previous algorithms [26, 30, 28] are based on applying standard (i.e., independent across jobs) randomized rounding to the fractional solution to find an assignment of jobs to machines. However, a very simple example (in Section 2) shows that no

such "*independent* randomized rounding" based algorithm can give a $3/2 - \Omega(1)$ guarantee, irrespective of the underlying convex relaxation. The problem is that the variance can be too high.

To get around this, we need to introduce some *strong* negative correlation among pairs of jobs assigned to any machine i (i.e., the ratio of the probability that they are both scheduled on i to the product of their respective probabilities of assigment on i, should be $1 - \Omega(1)$). For intuition, consider the example depicted on the left in Figure 1: we have a set $\{1, 2, 3, 4\}$ of four jobs, two machines a, b, and the SDP assigns each job fractionally $1/2$ to both machines. Note that independent randomized rounding would assign any two jobs j and j' to machine a (and similarly to b) with probability $1/4$. Ideally, we would like to have strong negative correlation that decreases this probability for all pairs of jobs. Unfortunately, this is not possible in general as can be seen by taking $n \gg 1$ jobs instead of 4 in the considered example[1]. However, one can still hope for a randomized rounding with strong negative correlation for some of the jobs while maintaining that no two jobs are assigned to a single machine with probability more than $1/4$. This is what our randomized rounding algorithm achieves.

The pairs of jobs that will have strong negative correlation are decided by a grouping scheme: for each machine i, the jobs are partitioned into groups with total fractional assignment on i being at most 1. The jobs in the same group are those that will have strong negative correlation. This step is illustrated on the right of Figure 1. Machine a has two groups consisting of jobs $\{1, 3\}$ and $\{2, 4\}$ and machine b has two groups consisting of jobs $\{1, 2\}$ and $\{3, 4\}$. Viewing this as a bipartite graph with group and job vertices, we would like to find an assignment with strong negative correlation on the edges incident to the same group. This is reminiscent of the several randomized pipage-based schemes [2, 3, 12, 14, 19] that given a fractional matching produce an integral matching. In fact, these get perfect negative correlation between edges at a vertex as only one edge is picked at any vertex[2].

However, these techniques do not work in our setting of general assignments due to a somewhat subtle issue; trying to force strong negative correlation between two edges in a group of a machine can cause unexpected positive correlations among other edges of that machine. In particular, in our example, previous rounding techniques would output one of the two perfect matchings with equal probability (the two perfect matchings are indicated by dashed and solid edges in the right of Figure 1) – thus yielding perfect *positive* correlation, e.g., for jobs 1 and 4 being assigned to machine a.

To get around this we give a new rounding theorem. The main idea behind the algorithm is to update the fractional assignment using randomized pipage steps along carefully chosen paths of length 4. In particular, these paths are chosen based on a random 2-coloring of the edges where the coloring is based on the fractional assignment and evolves over time. The properties of this general rounding technique are summarized in the theorem below. We believe that this technique can be of independent interest, as it appears to be the first to obtain *strong* negative correlations. Indeed, our rounding maintains the desired properties from independent randomized rounding (properties (a), (b), and the second part of (c) of Theorem 1.2) while also achieving a guaranteed amount of pairwise negative – i.e., strong – correlation (the first part of (c)). This is key for our result, and we are not aware of any prior work in this vein.

Theorem 1.2. *Let $\zeta = 1/108$. Consider a bipartite graph $G = (U \cup V, E)$ and let $y \in [0, 1]^E$ be fractional values on the edges satisfying $y(\delta(v)) = 1$ for all $v \in V$. For each vertex $u \in U$, select any family of disjoint $E_u^{(1)}, E_u^{(2)}, \ldots, E_u^{(\kappa_u)} \subseteq \delta(u)$ subsets of edges incident to u such that $y(E_u^{(\ell)}) \leq 1$ for $\ell = 1, \ldots, \kappa_u$. Then*

[1] Any schedule π of n jobs on 2 machines, must have $\Pr_{j,j'}[\,\pi$ assigns j and j' to the same machine$] \geq 1/2 - o(1)$. A simple proof of this is as follows. Suppose π assigns s jobs to the first machine and t jobs to the second machine, where $s + t = n$. Then, conditional on this, the desired probability is $(\binom{s}{2} + \binom{t}{2})/\binom{n}{2}$ which is minimized at $s = t = n/2$, and has value $1/2 - o(1)$. Now as $\mathbb{E}_\pi \Pr_{j,j'}[\,\pi$ assigns j and j' to the same machine$] = \mathbb{E}_{j,j'} \Pr_\pi[\,\pi$ assigns j and j' to the same machine$]$, we have that there exist two jobs j and j' that are assigned to the same machine with probability at least $1/2 - o(1)$ (and thus to one of them with probability at least $1/4 - o(1)$) no matter which algorithm, i.e., distribution over schedules π, that is used.

[2] Although they can introduce positive correlation between non-adjacent edges.

there exists a randomized polynomial-time algorithm that outputs a random subset of the edges $E^* \subseteq E$ satisfying

(a) For every $v \in V$, we have $|E^* \cap \delta(v)| = 1$ with probability 1;

(b) For every $e \in E$, $\Pr[e \in E^*] = y_e$;

(c) For every $u \in U$ and all $e \neq e' \in \delta(u)$,

$$\Pr[e \in E^* \wedge e' \in E^*] \leq \begin{cases} (1-\zeta) \cdot y_e y_{e'} & \text{if } e, e' \in E_u^{(\ell)} \text{ for some } \ell \in \{1, 2, \ldots, \kappa_u\}, \\ y_e y_{e'} & \text{otherwise}. \end{cases}$$

In the above theorem, we use the standard notation $\delta(w) = \{e \in E : w \in e\}$ to denote the set of edges incident to a vertex w, and let $y(F) = \sum_{e \in F} y_e$ for any subset $F \subseteq E$ of edges.

2 Preliminaries and Lower Bounds

On a single machine, the weighted completion is minimized by ordering the jobs in non-increasing order of w_j/p_j, referred to as the Smith ordering. In the unrelated machines setting, for each machine i let $\preceq_i$ denote the Smith ordering of jobs on machine i (i.e. $j' \preceq_i j$ iff $w_{j'}/p_{ij'} \geq w_j/p_{ij}$). Given an assignment of jobs to machines, the total weighted completion time is simply

$$\sum_i \sum_{j \in J(i)} w_j p_{ij} \Big(\sum_{j' \preceq_i j} p_{ij'} \Big)$$

where $J(i)$ denotes the set of jobs assigned to machine i.

For each $i \in M$ and $j \in J$, consider a binary variable x_{ij} that should take value 1 if and only if job j is assigned to machine i. Then the exact quadratic program can be formulated as follows:

$$\begin{aligned} \text{(QP)} \qquad \text{Minimize} \quad & \sum_{i \in M} \sum_{j \in J} w_j x_{ij} \left(\sum_{j' \in J : j' \preceq_i j} p_{ij'} x_{ij'} \right) \\ \text{subject to} \quad & \sum_{i \in M} x_{ij} = 1 \qquad \text{for all } j \in J, \\ & x \in \{0,1\}^{M \times N}. \end{aligned}$$

The Convex Programming relaxation of [30, 28]: We only describe the relaxation of [30, 28] here and refer to [30] for details on how it is obtained. They relax the variables x_{ij} in (QP) above to be fractional in $[0, 1]$, together with the fact that $x_{ij}^2 = x_{ij}$ for an integral solution and that $c^T x := \sum_i \sum_j w_j p_{ij} x_{ij}$ is a lower bound on any solution to obtain the following convex relaxation:

$$\begin{aligned} \text{(CP)} \qquad \text{Minimize} \quad & z \\ \text{subject to} \quad & z \geq \frac{1}{2} c^T x + \frac{1}{2} x^T D x \\ & z \geq c^T x \\ & \sum_{i \in M} x_{ij} = 1 \qquad \text{for all } j \in J, \\ & x \in [0,1]^{M \times N}. \end{aligned}$$

where $x^T Dx := \sum_i (\sum_j w_j (\sum_{j' \prec_{ij}} 2p_{ij'} x_{ij'} + p_{ij} x_{ij}) x_{ij})$ can be shown to be a convex function. We will refer to $c^T x$ and $x^T Dx$ as the linear and quadratic terms respectively.

A 3/2 integrality gap instance for CP: Consider the following instance. There are $k+1$ jobs, all of weight 1. The first k jobs are of size (processing time) 1 each and can only be placed on machine 1 (i.e., have infinite size on other machines). Job $k+1$ has size k^2 and can be placed on any machine $2, \ldots, m$, where we let $m = k+1$.

Claim 2.1. *The above instance has an integrality gap $3/2 - O(1/k)$ for (CP).*

Proof. First observe that any integral solution has value greater than $(3/2)k^2$ as the total completion time of the first k jobs is $k(k+1)/2$ while the last job has a completion time of k^2.

Now, consider the fractional solution where each job $1, \ldots, k$ is assigned to extent 1 on machine 1, and job $k+1$ is assigned to each machine i for $i = 2, \ldots, m$, to an extent of $1/(m-1) = 1/k$. We will show that this solution has fractional value at most $k^2 + k$.

First, the linear term $c^T x$ is $k + k^2$ (k for the first k jobs and k^2 for the big job). Second, the quadratic term is

$$x^T Dx = \sum_{j=1}^{k} (2(j-1)+1) + \sum_{i=2}^{m} \frac{k^2}{(m-1)^2} = k^2 + \frac{k^2}{m-1} = k^2 + k.$$

In particular, the first k jobs contribute k^2 above, and for the last job each of the $m-1$ machines contributes $k^2/(m-1)^2$. Thus (CP) has objective value at most $k^2 + k$. □

Note that in this example, the problem is that both the linear and quadratic bounds are weak on the overall instance. In particular, while the linear bound is exact on the big job, it is very weak on the small jobs. On the other hand, the quadratic term is exact on the small jobs, but very weak on the big job.

Limitation of Independent Randomized Rounding based approaches: The previous-best approximation algorithms are based on standard (i.e., *independent* across jobs) randomized rounding. We show that no such rounding can beat the approximation guarantee of $3/2$, irrespective of the relaxation. Consider the (trivial) instance with m jobs each of which can be placed on any of the m machines, and with $w_j = p_{ij} = 1$ for all i, j. The fractional solution $x_{ij} = 1/m$ for all $i, j \in [m]$ is a valid solution for any relaxation (as it is can be expressed as a convex combination of m perfect matchings). Clearly, the optimal solution assigns one job to each machine and has value m. However, under independent randomized rounding, for large m, the number of jobs assigned to a machine approaches a Poisson distribution with mean 1 and so the probability that a machine gets k jobs is $\approx 1/(e \cdot k!)$. The expected completion time on any machine is thus

$$\approx \sum_{k=0}^{\infty} \frac{k(k+1)}{2} \cdot \frac{1}{ek!}$$

which is $3/2$ as the first and second moments of Poisson(1) are 1 and 2 respectively.

The need for negative correlation in different classes: The above example might suggest that randomized rounding performs poorly only when the total mass ($\sum_j x_{ij}$) on a machine i is close to 1, as intuitively the effect of the variance should be relatively small if there are many jobs. This intuition is indeed true if the jobs are similar to each other in terms of size (processing time) and weight. However, the following example shows that some more care is needed if the jobs are very dissimilar. Suppose there are ℓ job classes $k = 1, \ldots, \ell$, where a class k job has weight M^k and size M^{-k} for some large M, and that machine i has m jobs from each class, with $x_{ij} = 1/m$ for all jobs j. So the total fractional assignment of jobs to i is ℓ. Now, as the Smith ratios are very different, the jobs from different classes have negligible effect on each other: only the individual cost of each class matters, and the fractional cost is $\approx \sum_{k=1}^{\ell} M^k M^{-k} = \ell$.

Now, if we round each job independently, the expected cost is $3\ell/2$, and it is not hard to see that to get a $((3/2) - c)$–approximation, we need to get a non-trivial negative correlation in at least $\Omega(c)$ fraction of the classes.

It turns out that this example is in a sense the worst possible; it motivates our rounding procedure in Section 5. Roughly speaking, it suffices to partition the jobs in different classes so that the total fractional weight is about 1, and then try to get some strong negative correlation within jobs of each class.

3 Strong Convex Relaxation

In this section, we give a strong convex relaxation based on the paradigm of "systematically" relaxing the exact quadratic mathematical program (QP) to a tractable convex program. In particular, our relaxation can be obtained "automatically" using the Lasserre/Sum-of-Squares hierarchy (although we have chosen to write this section in a self-contained manner).

To obtain a convex relaxation of (QP), we linearize it by replacing each quadratic term $x_{ij} \cdot x_{ij'}$ by a new variable $x_{\{ij,ij'\}}$ with the exception that $x_{ij} \cdot x_{ij}$ is replaced by the existing variable x_{ij} (since in any binary solution $x_{ij}^2 = x_{ij}$). For notational convenience, we also refer to variable x_{ij} as $x_{\{ij\}}$ and we introduce an auxiliary variable $x_\emptyset$ and set $x_\emptyset = 1$. The set of variables of our relaxation is thus $\{x_\emptyset\} \cup \{x_{\{ij\}\cup\{ij'\}}\}_{i\in M, j,j'\in J}$. Clearly any intended solution satisfies that $\sum_{i\in M} x_{ij} = 1$ and that x is non-negative. Another family of valid constraints is as follows. For a machine $i \in M$, let $X^{(i)}$ be the $(n+1)\times(n+1)$ matrix whose rows and columns are indexed by $\emptyset$ and $\{ij\}_{j\in J}$. The entries of $X^{(i)}$ are defined by $X^{(i)}_{S,T} = x_{S\cup T}$. In particular, this implies that $X^{(i)}_{\emptyset,\{ij\}} = X^{(i)}_{\{ij\},\{ij\}} = x_{ij}$ (that we will use crucially). We impose the constraint that $X^{(i)} \succeq 0$. These are valid constraints: indeed, if $X^{(i)}$ corresponds to an integral assignment x then

$$X^{(i)} = zz^T \succeq 0 \qquad \text{where} \quad z = (1, x_{i1}, \cdots, x_{in})^T$$

and $X^{(i)}_{\{ij\},\{ij\}} = (zz^T)_{ij,ij} = x_{ij}x_{ij} = x_{ij} = X^{(i)}_{\emptyset,\{ij\}}$.

The above yields the following convex (semidefinite programming) relaxation of our problem:

$$\begin{aligned}
\text{(SDP)} \quad \text{minimize} \quad & \sum_{i\in M}\sum_{j\in J} w_j \left(\sum_{j'\in J: j' \preceq_i j} p_{ij'} x_{\{ij\}\cup\{ij'\}} \right) \\
\text{subject to} \quad & \sum_{i\in M} x_{ij} = 1 \qquad \text{for all } j \in J, \\
& X^{(i)} \succeq 0 \qquad \text{for all } i \in M, \\
& x_\emptyset = 1, \\
& X^{(i)}_{S,T} \geq 0 \qquad \text{for all } i \in M \text{ and } S, T \subset J \text{ with } |S|, |T| \leq 1.
\end{aligned}$$

3.1 Lower bounds on the objective value

We briefly sketch why this SDP is stronger; e.g., it is exact on the $3/2$ integrality gap instance from Section 2.

Similar to previous works, our analysis reduces to that of fixing a single machine i and analyzing the cost of that machine: we compare the contribution of that machine to the objective of (SDP) to the (expected) cost of that machine in the schedule returned by our (randomized) algorithm. To do so, it will be important

to understand machine i's contribution to the objective when a job's processing time equals its weight, i.e., $p_{ij} = w_j$ for $j \in J$. In this case,

$$\sum_{j \in J} w_j \left(\sum_{j' \in J : j' \preceq_i j} p_{ij'} x_{\{ij\} \cup \{ij'\}} \right) = \sum_{j=1}^{n} p_{ij} (p_{i1} x_{\{ij\} \cup \{i1\}} + \cdots + p_{ij} x_{\{ij\} \cup \{ij\}}),$$

where we numbered the jobs according to the Smith ordering on machine i.

Interestingly, we can lower-bound this quantity is various ways as shown in the following lemma. The proof of this lemma crucially uses the SDP constraints and is deferred to the analysis of our approximation guarantee (see Lemma 5.4).

Lemma 3.1. *For any subset $S \subseteq \{1, \ldots, n\}$ of jobs,*

$$\sum_{j=1}^{n} p_{ij} (p_{i1} x_{\{ij\} \cup \{i1\}} + \cdots + p_{ij} x_{\{ij\} \cup \{ij\}}) \geq \sum_{j \notin S} x_{ij} p_{ij}^2 + \frac{1}{2} \left(\sum_{j \in S} x_{ij} p_{ij}^2 + \left(\sum_{j \in S} x_{ij} p_{ij} \right)^2 \right).$$

In particular, we can choose the best set S that gives us the tightest combination of the linear and the quadratic lower bounds. In contrast, the relaxations used in [30, 28] basically take the maximum lower bound (averaged over the machines) obtained by either setting $S = \emptyset$ or $S = J$.

This flexibility in choosing S will be critical to our analysis. For the $3/2$ gap instance, recall that the linear bound was tight for the large job, while the quadratic bound was tight for the small jobs, which makes the SDP exact on that instance.

4 Bipartite Assignment with Strong Negative Correlation

As discussed in Section 2, independent randomized rounding cannot give a better approximation ratio than $3/2$. To improve upon this ratio, we would ideally like to introduce strong negative correlation on jobs being assigned to a machine of the following type: if a job j is assigned to a machine, it should be less likely to assign other jobs to that machine. While it is not always impossible to introduce such negative correlations among all jobs, Theorem 1.2, which we prove in this section, shows that it is possible to introduce strong negative correlation between subsets of jobs (or vertices) without introducing positive correlations at pairs of edges with a common end-point. For convenience, we restate the theorem here.

Theorem 1.2. *Let $\zeta = 1/108$. Consider a bipartite graph $G = (U \cup V, E)$ and let $y \in [0,1]^E$ be fractional values on the edges satisfying $y(\delta(v)) = 1$ for all $v \in V$. For each vertex $u \in U$, select any family of disjoint $E_u^{(1)}, E_u^{(2)}, \ldots, E_u^{(\kappa_u)} \subseteq \delta(u)$ subsets of edges incident to u such that $y(E_u^{(\ell)}) \leq 1$ for $\ell = 1, \ldots, \kappa_u$. Then there exists a randomized polynomial-time algorithm that outputs a random subset of the edges $E^* \subseteq E$ satisfying*

(a) For every $v \in V$, we have $|E^* \cap \delta(v)| = 1$ with probability 1;

(b) For every $e \in E$, $\Pr[e \in E^*] = y_e$;

(c) For every $u \in U$ and all $e \neq e' \in \delta(u)$,

$$\Pr[e \in E^* \wedge e' \in E^*] \leq \begin{cases} (1-\zeta) \cdot y_e y_{e'} & \text{if } e, e' \in E_u^{(\ell)} \text{ for some } \ell \in \{1, 2, \ldots, \kappa_u\}, \\ y_e y_{e'} & \text{otherwise}. \end{cases}$$

We start by describing the randomized algorithm and then give its analysis.

Notation: Floating values. A value $z \in [0,1]$ will be called "floating" if $z \in (0,1)$.

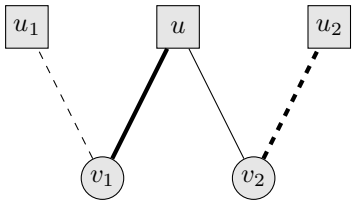

Figure 2: Illustration of the update in phase 2. Solid edges are in R and either (i) thick edges are increased by α and slim edges are decreased by α or (ii) slim edges are increased by β and thick edges are decreased by β. We note that u_1 may equal u_2 but they both differ from u.

4.1 Algorithm

We divide the algorithm into three phases and present each phase along with some simple observations that will be useful in the analysis.

Phase 1 (Forming the collection R^*): Let y^* denote the initial fractional assignment. For each vertex $v \in V$, partition its incident edges $\delta(v)$ into at most 6 disjoint groups by letting each group –except possibly for at most one group – be a minimal set of incident edges whose y^*-values sum up to at least $1/6$. (Note that this results in at most 6 groups since $y^*(\delta(v)) = 1$, and that these groups can be formed arbitrarily by picking the edges in $\delta(v)$ greedily in non-increasing order of y^*-value; the last group may have y^*-value smaller than $1/6$.) Now select a random group, uniformly at random and independently for each vertex v, and let R^* be the set of selected edges.

Observation 4.1. *Let $e, e' \in \delta(u)$ for some $u \in U$. Then,* $\Pr[(e \in R^*) \wedge (e' \in R^*)] \geq 1/36$.

Proof. The events that $e \in R^*$ and that $e' \in R^*$ are independent as they both are incident to different vertices in V. Now the statement follows as each $v \in V$ selects a random group out of at most 6 many. □

Phase 2 (Updating the assignment): Initially let $y = y^*, R = R^*$. Repeat the following steps while there exist edges $\{u, v_1\}, \{u, v_2\} \in R \cap E_u^{(\ell)}$ for some ℓ and $\{u_1, v_1\} \in \delta(v_1) \setminus R$ and $\{u_2, v_2\} \in \delta(v_2) \setminus R$ with *floating* y-value. Here $u, u_1, u_2 \in U$, $v_1, v_2 \in V$, but are otherwise arbitrary. See Figure 2:

1. Let $\alpha = \min\{y_{u_1,v_1}, 1 - y_{u,v_1}, y_{u,v_2}, 1 - y_{u_2,v_2}\}$ and $\beta = \min\{1 - y_{u_1,v_1}, y_{u,v_1}, 1 - y_{u,v_2}, y_{u_2,v_2}\}$.
2. With probability $\frac{\alpha}{\beta+\alpha}$, update y as follows for each $e \in E$:

$$y_e = \begin{cases} y_e + \beta & \text{if } e = \{u_1, v_1\} \text{ or } e = \{u, v_2\}, \\ y_e - \beta & \text{if } e = \{u, v_1\} \text{ or } e = \{u_2, v_2\}, \\ y_e & \text{otherwise.} \end{cases}$$

Otherwise (with remaining probability $\frac{\beta}{\alpha+\beta}$), update y as follows for each $e \in E$:

$$y_e = \begin{cases} y_e - \alpha & \text{if } e = \{u_1, v_1\} \text{ or } e = \{u, v_2\}, \\ y_e + \alpha & \text{if } e = \{u, v_1\} \text{ or } e = \{u_2, v_2\}, \\ y_e & \text{otherwise.} \end{cases}$$

3. For $v \in \{v_1, v_2\}$, if $\sum_{e \in \delta(v) \cap R} y_e = 1$, i.e. if all the edges incident to v are in R, then update R as

$$R = (R \setminus \delta(v)) \cup \left\{ \arg \max_{e \in \delta(v) \cap R} y_e \right\}.$$

That is, remove all edges incident to v from R, except one with the largest y-value.

We note the following simple observations about this phase.

Observation 4.2. *During Phase 2, if a variable y_e reaches 0 or 1, then it is not updated anymore. Moreover, at each iteration of Phase 2, at least one edge with floating y-value has its y-value reach 0 or 1.*

Proof. This follows from that Phase 2 only updates floating y-values and, in each iteration, α and β is selected so that one of the selected edges' y-value will reach 0 or 1. □

Observation 4.3. *Phase 2 satisfies the invariants $y(\delta(v)) = 1$ for every $v \in V$ and $y_e \geq 0$ for every $e \in E$.*

Proof. Notice that when y is updated then the selection of α and β guarantees that $y_e \geq 0$ for every $e \in E$. Moreover, the update is designed so that the fractional degree of a vertex in V stays constant. Thus the statement follows since we start with $y = y^*$ for which $y(\delta(v)) = 1$ for $v \in V$. □

Observation 4.4. *The set R does not increase in size during Phase 2. Moreover, if an edge $e \in \delta(v) \cap R$ is removed from R (in Step 3) then it must be that $y(e) \leq 1/2$ after Step 2.*

Proof. That R only decreases in size follows directly from Step 3. For the second part, if Step 3 is applied at v and $y(e) > 1/2$ for some $e \in \delta(v)$, then as $\sum_{e' \in \delta(v)} y(e') = 1$, it must be that $e = \arg\max_{e' \in \delta(v)}$ and thus e remains in R. □

Observation 4.5. *When Phase 2 terminates, then for every $u \in U$ and $\ell \in \{1, \ldots, \kappa_u\}$, we have $|\{e \in E_u^{(\ell)} \cap R \mid y_e > 0\}| \leq 1$.*

Proof. Suppose that there exist $e_1, e_2 \in E_u^{(\ell)} \cap R$ with $y_{e_1}, y_{e_2} > 0$. Then since any iteration of Phase 2 maintains the value of $y(E_u^{(\ell)} \cap R)$ and $R \subseteq R^*$ we have $y(E_u^{(\ell)} \cap R) \leq y^*(E_u^{(\ell)} \cap R^*) \leq 1$. Hence, $y_{e_1}, y_{e_2} < 1$. Now by Step 3 of Phase 2, we are guaranteed that a not-yet-integrally-assigned vertex $v \in V$ has $y(\delta(v) \cap R) < 1$. Therefore, there exist edges $e_1 = \{v_1, u\}, e_2 = \{v_2, u\}$ and $\{u_1, v_1\} \in \delta(v_1) \setminus R$ and $\{u_2, v_2\} \in \delta(v_2) \setminus R$ with floating y-values. This implies that Phase 2 does not terminate in this case. □

Phase 3 (Randomized Rounding): Form E^* by, independently for each vertex $v \in V$, selecting a single edge $e \in \delta(v)$ so that $e \in \delta(v)$ is selected with probability y_e. Notice that this is possible because, by Observation 4.3, we have $\sum_{e \in \delta(v)} y_e = 1$ for all $v \in V$ and $y_e \geq 0$ for all $e \in E$.

4.2 Analysis

We first note that the algorithm terminates in polynomial time. Phase 1 and Phase 3 both clearly run in polynomial time. Each step of Phase 2 runs in polynomial time and by Observation 4.2, Phase 2 runs in at most $|E|$ iterations.

We continue to analyze the properties. The intuition for why they should hold is as follows. The algorithm is inspired by randomized-rounding algorithms for bipartite matchings such as pipage rounding and swap rounding. It is easy to see that these algorithms satisfy both Property (a) and the marginal probabilities (Property (b)): indeed, α and β are defined in order to do so. Moreover, the weak bound of Property (c) follows basically from the fact that, for each $u \in U$, the y-values of two edges incident to u are never increased

simultaneously. Finally, the intuition behind the novel strong bound of Property (c) is as follows. After Phase 2, the probability that two vertices $e, e' \in E_u^{(\ell)}$ are in R is at least $1/36$. Now using that the initial y-value of edges in $\delta(v) \cap R$ is at most $1/3$ for every $v \in V$, and that the y-values of edges are preserved in expectation, there is a reasonable probability that both e, e' will remain in R until the end. However, in that case, it is easy to see by Observation 4.5 that at most one of them will be selected in E^*. We now continue to formally prove these properties.

Property (a)**:** That Property (a) of Theorem 1.2 holds follows from Observation 4.3 and as Phase 3 chooses exactly one edge incident to each $v \in V$.

Properties (b) **and** (c)**:** To show these properties, we will inductively show some invariants. Let $Y^{(k)} = (y_e^{(k)} : e \in E)$ denote the collection of y-values of edges and $R^{(k)}$ be the set R at the end of iteration k of Phase 2. For an edge $e = \{u, v\} \in R$ with $u \in U$ and $v \in V$ let $R_{\bar{e}} = \{e' \in \delta(v) \cap R : e' \neq e\}$ be the other edges in R incident to v.

We show the following invariants hold after each iteration k. Here, conditioning an event on $Y^{(k)}$ and $R^{(k)}$ means the probability of that event if the random iterations in Phase 2 are applied starting from the assignment $Y^{(k)}$ and $R = R^{(k)}$.

$$\Pr[e \in E^* \mid Y^{(k)}, R^{(k)}] = y_e^{(k)} \qquad \forall e \in E \tag{1}$$

$$\Pr[e \in E^* \wedge e' \in E^* \mid Y^{(k)}, R^{(k)}] \leq y_e^{(k)} y_{e'}^{(k)} \qquad \forall u \in U, e, e' \in \delta(u) \tag{2}$$

$$\Pr[e \in E^* \wedge e' \in E^* \mid Y^{(k)}, R^{(k)}] \leq 2(y^{(k)}(R_{\bar{e}}^{(k)}) + y^{(k)}(R_{\bar{e}'}^{(k)}))y_e^{(k)} y_{e'}^{(k)} \tag{3}$$
$$\forall u \in U, \ell \in \{1, \ldots, \kappa_u\}, e \neq e' \in E_u^{(\ell)} \cap R^{(k)}$$

To show these, we will apply reverse induction. For the base case, we show that these properties hold after the last iteration of Phase 2. For the inductive step, we show that if they hold after the k-th iteration then they also hold after iteration $k-1$ (or equivalently at the beginning of iteration k), and hence they also hold for the y-values and R at the beginning of Phase 2.

Let us first see how this implies the theorem.

At the beginning of Phase 2 we have $y^{(0)} = y^*$ and $R^{(0)} = R^*$. So having (1) for $k = 0$, implies Property (b) and (2) implies the weaker bound in Property (c). For the stronger bound, consider two edges $e \neq e' \in E_u^{(\ell)}$. By (3) above, we have that

$$\begin{aligned}
\Pr[e \in E^* \wedge e' \in E^*] &= \mathbb{E}_{R^*}[\Pr[(e \in E^* \wedge e' \in E^*) \mid Y^*, R^*]] \\
&\leq \Pr[e, e' \in R^*] \cdot 2(y^*(R_{\bar{e}}) + y^*(R_{\bar{e}'}))y_e^* y_{e'}^* + (1 - \Pr[e, e' \in R^*]) \cdot y_e^* y_{e'}^* \\
&\leq \Pr[e, e' \in R^*] \frac{2y_e^* y_{e'}^*}{3} + (1 - \Pr[e, e' \in R^*])y_e^* y_{e'}^* \\
&\leq \frac{2y_e^* y_{e'}^*}{3 \cdot 36} + \frac{35}{36} y_e^* y_{e'}^* \\
&= \frac{107}{108} y_e^* y_{e'}^*.
\end{aligned}$$

The second inequality follows from the fact that $y^*(R_{\bar{e}}), y^*(R_{\bar{e}'}) \leq 1/6$ because $e, e' \in R$ and after Phase 1 $R \cap \delta(v)$ is a *minimal* group with y^*-value at least $1/6$ for each $v \in V$; and the third inequality follows from Observation 4.1.

It thus remains to prove (1)-(3) by reverse induction on the iterations in Phase 2. One subtle point in the argument is that the set R might also change (reduce in size) after an iteration.

Base case (when Phase 2 terminates): In this case Phase 2 will not change any of the y-values. As each vertex $v \in V$ picks an edge in $\delta(v)$ randomly with probability y_e, $\Pr[e \in E^*] = y_e$ for every

$e \in E$, so (1) is satisfied. Similarly for (2), we note that for two edges $e \neq e' \in \delta(u)$, it holds that $\Pr[e \in E^* \wedge e' \in E^*] = y_e y_{e'}$.

Finally, Observation 4.5 says that the number of edges in $E_u^{(\ell)} \cap R$ with positive y-value is at most 1. Therefore, we have that $\Pr[e \in E^* \wedge e' \in E^*] = 0$ for $e \neq e' \in E_u^{(\ell)} \cap R$ and (3) holds trivially.

Inductive step: Assuming (1)-(3) holds at the end of iteration k, we prove that they hold at the end of iteration $k-1$).

For notational ease, let us denote $Y = Y^{(k-1)}$, $R = R^{(k-1)}$ and let $Y' = Y^k$ and $R' = R^{(k)}$ denote the (random) updated y-values and set R.

We first verify (1). By the inductive hypothesis (I.H.) we have that $\Pr[e \in E^* \mid Y'] = y'_e$. So, $\Pr[e \in E^* \mid Y] = \mathbb{E}_{Y'|Y}[\Pr[e \in E^* \mid Y']]$, which is $\mathbb{E}_{Y'|Y}[y'_e]$. If Phase 2 did not update the value of edge e then clearly $y'_e = y_e$. Otherwise, we have $\mathbb{E}_{Y'|Y}[y'_e] = \frac{\alpha}{\alpha+\beta}(y_e + \beta) + \frac{\beta}{\alpha+\beta}(y_e - \alpha) = y_e$. Thus, (1) holds in either case.

Similarly, we show (2). By the I.H., $\Pr[e \in E^* \wedge e' \in E^* | Y'] \leq y'_e y'_{e'}$ and thus

$$\Pr[e \in E^* \wedge e' \in E^* \mid Y] = \mathbb{E}_{Y'|Y}[\Pr[e \in E^* \wedge e' \in E^* \mid Y'] \leq \mathbb{E}_{Y'|Y}[y'_e y'_{e'}].$$

On the one hand, if Phase 2 only changed the y-value for at most one of e and e', then by independence this is at most $\mathbb{E}_{Y'|Y}[y'_e]\mathbb{E}_{Y'|Y}[y'_{e'}] = y_e y_{e'}$. On the other hand, if it changed both of the values then we have

$$\mathbb{E}_{Y'|Y}[y'_e y'_{e'}] = \frac{\alpha}{\alpha+\beta}(y_e + \beta)(y_{e'} - \beta) + \frac{\beta}{\alpha+\beta}(y_e - \alpha)(y_{e'} + \alpha) \leq y_e y_{e'}.$$

Indeed, if Phase 2 changes the value of two edges incident to a vertex in U then it always increases the value of one edge and decreases the value of the other edge. We have thus that (2) is satisfied.

We finish the analysis by verifying (3). Consider $e \neq e' \in E_u^{(\ell)} \cap R$ for some $u \in U$ and $\ell \in \{1, \ldots, \kappa_u\}$. We wish to show that

$$\Pr[e \in E^* \wedge e' \in E^* \mid Y] \leq 2(y(R_{\bar{e}}) + y(R_{\bar{e}'}))y_e y_{e'}.$$

Let R' be the set R after the single iteration of Phase 2. As previously we will use that

$$\Pr[e \in E^* \wedge e' \in E^* \mid Y] = \mathbb{E}_{Y'|Y}[\Pr[e \in E^* \wedge e' \in E^* \mid Y']$$

and the I.H., but we cannot do it directly as it might the case that even though e and e' belong to R, they may not belong to R'. So we condition the right hand side depending on whether this happens or not.

Suppose $e \notin R'$. Then by Observation 4.4, $y'(R_{\bar{e}}) \geq 1/2$ and hence we have that $2(y'(R_{\bar{e}}) + y'(R_{\bar{e}'}))y'_e y'_{e'} \geq y'_e y'_{e'}$. By (2) we have that

$$\Pr[e \in E^* \wedge e' \in E^* \mid Y', R'] \leq y'_e y'_{e'},$$

this implies that (conditioned on $e \notin R'$)

$$\Pr[e \in E^* \wedge e' \in E^* \mid Y', R'] \leq 2(y'(R_{\bar{e}}) + y'(R_{\bar{e}'}))y'_e y'_{e'}.$$

The same holds if $e' \notin R'$.

Now if both e and e' lie in R' by the I.H. we know that

$$\Pr[e \in E^* \wedge e' \in E^* \mid Y', R'] \leq 2(y'(R'_{\bar{e}}) + y'(R'_{\bar{e}'}))y'_e y'_{e'} \leq 2(y'(R_{\bar{e}}) + y'(R_{\bar{e}'}))y'_e y'_{e'},$$

where the last inequality follows from Observation 4.4, i.e., from the fact that $R' \subseteq R$.

We have thus upper bounded all cases (irrespective of whether R' contains e or e') by the same expression and it suffices to show

$$\mathbb{E}_{Y'|Y}[2(y'(R_{\bar{e}}) + y'(R_{\bar{e}'}))y'_e y'_{e'}] \leq 2(y(R_{\bar{e}}) + y(R_{\bar{e}'}))y_e y_{e'}$$

If neither e or e' is changed by the iteration of Phase 2, then

$$\mathbb{E}_{Y'|Y}[2(y'(R_{\bar{e}}) + y'(R_{\bar{e}'}))y'_e y'_{e'}] = \mathbb{E}_{Y'|Y}[2(y'(R_{\bar{e}}) + y'(R_{\bar{e}'}))]y_e y_{e'} = 2(y'(R_{\bar{e}}) + y'(R_{\bar{e}'}))y_e y_{e'},$$

where the second equality follows by linearity of expectation and (1).

Now suppose the iteration of Phase 2 changes at least one of e or e'. Then we claim that $y'(R_{\bar{e}}) = y(R_{\bar{e}})$ and $y'(R_{\bar{e}'}) = y(R_{\bar{e}'})$. To see this note that an iteration of Phase 2 changes exactly two edges in R incident to the same vertex in U and since, in this case, one of them is incident to u so must the other one. The sets $R_{\bar{e}}$ and $R_{\bar{e}'}$ only contain edges of R that are not incident to u and are thus left unchanged in this case. Hence,

$$\begin{aligned}\mathbb{E}_{Y'|Y}[2(y'(R_{\bar{e}}) + y'(R_{\bar{e}'}))y'_e y'_{e'}] &= 2(y(R_{\bar{e}}) + y(R_{\bar{e}'}))\mathbb{E}_{Y'|Y}[y'_e y'_{e'}] \\ &\leq 2(y(R_{\bar{e}}) + y(R_{\bar{e}'}))y_e y_{e'},\end{aligned}$$

where the last inequality follows from (2). We have thus also proved (3) which completes the proof.

5 Rounding the Fractional Schedule

We now describe our scheduling algorithm. The algorithm solves the SDP relaxation from Section 3, and applies the bipartite rounding procedure from Section 4 to a suitably defined graph based on the SDP solution. We will analyze this algorithm in Section 5.2, and in particular show the following result which directly implies Theorem 1.1.

Theorem 5.1. *The expected cost of the rounding algorithm is at most $(3/2 - c)$ times the cost of the optimal solution to the relaxation, where $c = \zeta/20000$ and ζ is the constant in Theorem 1.2.*

5.1 Description of Algorithm

Our rounding algorithm consists of defining groups (i.e., the families $E_u^{(\ell)}$) for each machine and then applying Theorem 1.2. Specifically, let x denote an optimal solution to our relaxation. We shall interpret the vector $y = (x_{ij})_{i \in M, j \in J}$ as an fractional assignment of jobs to machines in the bipartite graph $G = (M \cup J, E)$ where $E = \{ij : y_{ij} > 0\}$. Notice, that $y(\delta(j)) = 1$ for each $j \in J$ and $y \geq 0$. Thus, y satisfies the assumptions of Theorem 1.2. It remains to partition the edges incident to the machines into groups. To do this, we apply the following grouping procedure to each machine separately.

Grouping Procedure: For a fixed machine i we define the groups as follows:

1. Call a job j of class k, if $p_{ij} \in [10^{k-1}, 10^k)$. We assume (by scaling) that $p_{ij} \geq 1$ if $p_{ij} \neq 0$.
2. For each class $k = 0, 1, 2, \ldots$, order the jobs in that class in non-increasing order of Smith's ratio, i.e., w_j/p_{ij}, and form groups as follows. If some job j has $x_{ij} \geq 1/10$, it forms a separate group by itself $\{j\}$. For the remaining jobs, greedily pick the jobs in class k so that their total fractional y-value on i first reaches at least $1/10$ and make it a group; and repeat until the remaining jobs of that class have total fractional value less than $1/10$.

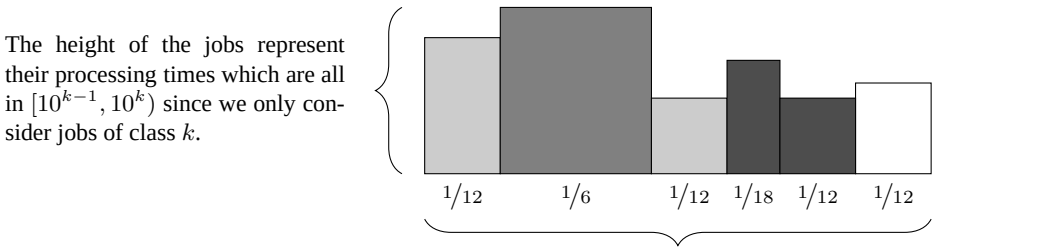
The height of the jobs represent their processing times which are all in $[10^{k-1}, 10^k)$ since we only consider jobs of class k.
$1/12$
$1/6$
$1/12$
$1/18$
$1/12$
$1/12$
The jobs are ordered in non-increasing order of Smith's ratio and the widths of the depicted jobs show their y-value on the considered machine.

Figure 3: Example of the grouping procedure on a machine i for the jobs of class k. The different groups are depicted in different colors; the job corresponding to the white rectangle is ungrouped.

By definition, the ungrouped jobs in each size class k have total fractional value less than $1/10$ on machine i. Note also that several singleton groups could be interspersed between jobs of a single group. For an example see Figure 3.

Let $E_i^{(1)}, \ldots, E_i^{(\kappa_i)}$ denote the groups formed, over all the classes, for machine i. We now apply Theorem 1.2 to the graph $G = (M \cup J, E)$ with $U = M$ and the groups $E_u^{(1)}, \ldots, E_u^{(\kappa(u))}$ at the machine $u \in U$. Observe that the conditions of the groups are satisfied, i.e., they are disjoint and the total y-value is at most 1 in all of them. This gives an assignment of the jobs to machines and thus a schedule.

5.2 Analysis

To analyze the performance of the algorithm above, we proceed in several steps. We first define some notation and make some observations that allow us to express the cost of the algorithm and the relaxation in a more convenient form. In Section 5.2.2 we show how to upper bound the cost of the schedule produced by the algorithm. In secton 5.2.3 we show how to derive various strong lower bounds from the SDP formulation, and finally in Section 5.2.4 we show how to combine these results to obtain Theorem 5.1.

5.2.1 Notation

Let X_{ij} denote the random indicator variable that takes value 1 if the algorithm assigns job j to machine i. The expected value of the returned schedule of the algorithm can then be written as $\sum_{i \in M} \mathrm{ALG}_i$, where ALG_i denotes the expected cost of machine i, i.e.,

$$\mathrm{ALG}_i = \mathbb{E}\left[\sum_{j \in J} X_{ij} w_j \left(\sum_{j' \preceq j} X_{ij'} p_{ij'}\right)\right] = \sum_{j \in J} w_j \left(\sum_{j' \preceq j} p_{ij'} \mathbb{E}[X_{ij} X_{ij'}]\right).$$

Similarly, the value of the optimal solution x to the relaxation can be decomposed into a sum $\sum_{i \in M} \mathrm{REL}_i$ over the costs of the machines, where

$$\mathrm{REL}_i = \sum_{j \in J} w_j \left(\sum_{j' \in J : j' \preceq j} p_{ij'} x_{\{ij\} \cup \{ij'\}}\right).$$

In order to prove Theorem 5.1, it is thus sufficient to show

$$\mathrm{ALG}_i \leq (3/2 - c)\,\mathrm{REL}_i \qquad \text{for all } i \in M. \tag{4}$$

To this end, we fix an arbitrary machine $i \in M$ and use the following notation:

- For simplicity, we abbreviate p_{ij} by p_j, x_{ij} by x_j, $x_{\{ij\}\cup\{ij'\}}$ by $x_{\{j\}\cup\{j'\}}$, and X_{ij} by X_j.
- We let $\beta_j = w_j/p_j$ denote Smith's ratio of job $j \in J$ on machine i and rename the jobs $J = \{1, 2, \ldots, n\}$ so that $\beta_1 \leq \beta_2 \leq \cdots \leq \beta_n$.

With this notation, we can rewrite REL_i and ALG_i as follows.

Lemma 5.2. *We have*

$$\textit{ALG}_i = \sum_{j=1}^{n} (\beta_j - \beta_{j+1}) \mathbb{E}\left[\sum_{j'=1}^{j} p_{j'} X_{j'} (p_1 X_1 + \cdots + p_{j'} X_{j'})\right],$$

$$\textit{REL}_i = \sum_{j=1}^{n} (\beta_j - \beta_{j+1}) \left[\sum_{j'=1}^{j} p_{j'} (p_1 x_{\{j'\}\cup\{1\}} + \cdots + p_{j'} x_{\{j'\}\cup\{j'\}})\right],$$

where for notational convenience we let $\beta_{n+1} = 0$.

Proof. We prove the first equality based on a telescoping sum argument. The second equality follows exactly by the same arguments. Using $w_j = \beta_j p_j$ we can rewrite

$$\text{ALG}_i = \mathbb{E}\left[\sum_{j=1}^{n} w_j \left(\sum_{j'=1}^{j} X_j X_{j'} p_{j'}\right)\right] = \mathbb{E}\left[\sum_{j=1}^{n} \beta_j p_j X_j \left(\sum_{j'=1}^{j} X_{j'} p_{j'}\right)\right].$$

We now claim that the right-hand side of this expression equals

$$\sum_{j=1}^{n} (\beta_j - \beta_{j+1}) \mathbb{E}\left[\sum_{j'=1}^{j} p_{j'} X_{j'} (p_1 X_1 + \cdots + p_{j'} X_{j'})\right].$$

Consider any term $p_k X_k p_\ell X_\ell$ with $k \leq \ell$. This term appears in $\mathbb{E}\left[\sum_{j=1}^{n} \beta_j p_j X_j \left(\sum_{j'=1}^{j} X_{j'} p_{j'}\right)\right]$ only when $j' = k$ and $j = \ell$ and has a coefficient of β_ℓ. The same term appears in the expression $\sum_{j=1}^{n} (\beta_j - \beta_{j+1}) \mathbb{E}\left[\sum_{j'=1}^{j} p_{j'} X_{j'} (p_1 X_1 + \cdots + p_{j'} X_{j'})\right]$ when $j = \ell, \ell+1, \ldots, n$ with coefficients $(\beta_\ell - \beta_{\ell+1}), (\beta_{\ell+1} - \beta_{\ell+2}), \ldots, (\beta_n - \beta_{n+1})$. Thus, by telescoping, the coefficient in front of $p_k X_k p_\ell X_\ell$ is again β_ℓ. ☐

By combining the above lemma with (4), we have further reduced our task of proving Theorem 5.1 to that of proving

$$\mathbb{E}\left[\sum_{j=1}^{n'} p_j X_j (p_1 X_1 + \cdots + p_j X_j)\right] \leq (3/2 - c) \left[\sum_{j=1}^{n'} p_j (p_1 x_{\{j\}\cup\{1\}} + \cdots + p_j x_{\{j\}\cup\{j\}})\right]. \quad (5)$$

for all $n' \in J$. The rest of this section is devoted to proving this inequality for a fixed n'. We shall use the following notation:

- Let G denote those jobs that are in the groups that only contain jobs from $\{1, \ldots, n'\}$. Let $\overline{G} = \{1, \ldots, n'\} \setminus G$ denote the "ungrouped" jobs. Note that, by the definition of the algorithm, specifically, the grouping, we have that each job class has fractional value less than $1/10$ in $\overline{G}$. Let $\mathcal{G}$ denote the collection of these groups restricted to jobs $\{1, \ldots, n'\}$.

- Let $L = \sum_{j=1}^{n'} x_j p_j$ denote the "linear" sum and let $Q = \sum_{j=1}^{n'} x_j p_j^2$ denote the "quadratic" sum. We also use the notation $\overline{L}$ and $\overline{Q}$ to denote the linear and quadratic sums when restricted to ungrouped jobs, i.e., $\overline{L} = \sum_{j \in \overline{G}} x_j p_j$ and $\overline{Q} = \sum_{j \in \overline{G}} x_j p_j^2$.

The proof of (5) is described over the following three subsections. In Section 5.2.2 we give an upper bound on the left-hand-side (LHS) of (5); in Section 5.2.3 we give several lower bounds on the right-hand-side (RHS) of (5); finally, in Section 5.2.4 we combine these bounds to prove (5).

5.2.2 Upper bound on the LHS of (5)

We give the following upper bound on the LHS of (5). The lemma essentially say that we have a "gain" of $O(\zeta)$ for each grouped job, which follows from our negative correlation rounding.

Lemma 5.3. *For $Q, \overline{Q}$ and L as defined above, we have*

$$\mathbb{E}\left[\sum_{j=1}^{n'} p_j X_j (p_1 X_1 + \cdots + p_j X_j)\right] \leq (1 - \zeta/200) \cdot Q + \zeta/200 \cdot \overline{Q} + 1/2 \cdot L^2.$$

Proof. Using that $X_j^2 = X_j$ and a simple recombination of the terms, we have that

$$\mathbb{E}\left[\sum_{j=1}^{n'} p_j X_j (p_1 X_1 + \cdots + p_j X_j)\right] = \mathbb{E}\left[\frac{1}{2}\sum_{j=1}^{n'} X_j p_j^2 + \frac{1}{2}\left(\sum_{j=1}^{n'} X_j p_j\right)^2\right]$$

As our rounding satisfies the marginals, this can be simplified to $\frac{1}{2}\sum_{j=1}^{n'} x_j p_j^2 + \frac{1}{2}\mathbb{E}\left[\left(\sum_{j=1}^{n'} X_j p_j\right)^2\right]$. We now upper bound the latter term.

$$\mathbb{E}\left[\left(\sum_{j=1}^{n'} X_j p_j\right)^2\right] = \mathbb{E}\left[\sum_{j,j'} X_j X_{j'} p_j p_{j'}\right] = \mathbb{E}\left[\sum_{j,j':j\neq j'} X_j X_{j'} p_j p_{j'}\right] + \mathbb{E}\left[\sum_{j:j=j'} X_j X_{j'} p_j p_{j'}\right]$$

$$\leq \left(\sum_{j\neq j'} x_j x_{j'} p_j p_{j'} - \zeta \sum_{G' \in \mathcal{G}} \sum_{j\neq j' \in G} x_j x_{j'} p_j p_{j'}\right) + \sum_j x_j p_j^2 \qquad \text{(by Theorem 1.2 and } \mathbb{E}[X_j^2] = \mathbb{E}[X_j] = x_j)$$

$$= \sum_{j,j'} x_j x_{j'} p_j p_{j'} + \sum_j (x_j - x_j^2) p_j^2 - \zeta \sum_{G' \in \mathcal{G}} \sum_{j\neq j' \in G'} x_j x_{j'} p_j p_{j'}$$

$$\leq \left(\sum_j x_j p_j\right)^2 + \sum_j x_j p_j^2 - \zeta \sum_{G' \in \mathcal{G}} \left(\sum_{j \in G'} x_j p_j\right)^2 \qquad \text{(since } \zeta \leq 1\text{)}.$$

Now, for each group $G' \in \mathcal{G}$, we have $\sum_{j \in G'} x_j \geq 1/10$ and $p_j \geq p_{j'}/10$ for $j, j' \in G$. Therefore,

$$\left(\sum_{j \in G'} x_j p_j\right)^2 \geq \sum_{j \in G'} x_j p_j^2 / 100.$$

Thus, we have that the expected cost of the machine is upper bounded by

$$\begin{aligned}
&\left(\sum_{j=1}^{n'} x_j p_j^2\right) + \frac{1}{2}\left(\left(\sum_{j=1}^{n'} x_j p_j\right)^2 - \left(\zeta \sum_{G' \in \mathcal{G}} \sum_{j \in G'} x_j p_j^2/100\right)\right) \\
&= \left(\sum_{j=1}^{n'} x_j p_j^2\right) + \frac{1}{2}\left(\left(\sum_{j=1}^{n'} x_j p_j\right)^2 - \left(\zeta \sum_{j \in G} x_j p_j^2/100\right)\right) \\
&= (1 - \zeta/200)\left(\sum_{j=1}^{n'} x_j p_j^2\right) + \zeta/200\left(\sum_{j \in \overline{G}} x_j p_j^2\right) + 1/2\left(\sum_{j=1}^{n'} x_j p_j\right)^2.
\end{aligned}$$

□

5.2.3 Lower bounds on the RHS of (5)

The lemma below gives a general lower bound that allows us to the RHS of (5) in various ways by choosing different subsets S. The particular lower bounds that we later use (by plugging particular choices of S) are then stated in Corollary 5.5.

Lemma 5.4. *For any subset $S \subseteq \{1, \ldots, n'\}$ of jobs*[3]*,*

$$\sum_{j=1}^{n'} p_j(p_1 x_{\{j\}\cup\{1\}} + \cdots + p_j x_{\{j\}\cup\{j\}}) \geq \sum_{j \notin S} x_j p_j^2 + \frac{1}{2}\left(\sum_{j \in S} x_j p_j^2 + \left(\sum_{j \in S} x_j p_j\right)^2\right).$$

Proof. Similar to the proof of Lemma 5.3,

$$\sum_{j=1}^{n'} p_j(p_1 x_{\{j\}\cup\{1\}} + \cdots + p_j x_{\{j\}\cup\{j\}}) = \frac{1}{2}\left(\sum_{j=1}^{n'} x_j p_j^2 + \sum_{j,j'=1}^{n'} x_{\{j\}\cup\{j'\}} p_j p_{j'}\right).$$

As x and p are non-negative vectors, ignoring the terms $x_{\{j\}\cup\{j'\}}$ with $j \in S$ and $j' \in S'$, this can be lower bounded by

$$\underbrace{\frac{1}{2}\left(\sum_{j \notin S} x_j p_j^2 + \sum_{j,j' \notin S} x_{\{j\}\cup\{j'\}} p_j p_{j'}\right)}_{(I)} + \underbrace{\frac{1}{2}\left(\sum_{j \in S} x_j p_j^2 + \sum_{j,j' \in S} x_{\{j\}\cup\{j'\}} p_j p_{j'}\right)}_{(II)}.$$

Again using that x and p are non-negative, ignoring the terms with $j \neq j$ we also have that

$$\left(\sum_{j,j' \notin S} p_j p_{j'} x_{\{j\}\cup\{j'\}}\right) \geq \sum_{j \notin S} x_j p_j^2.$$

Hence, $(I) \geq \sum_{j \notin S} x_j p_j^2$.

[3]Here, and in the following, we mean $j \in \{1, \ldots, n'\} \setminus S$ by $j \notin S$.

Let us now concentrate on (II) and in particular we show that

$$\sum_{j,j' \in S} x_{\{j\}\cup\{j'\}} p_j p_{j'} \geq \mu^2$$

where $\mu = \sum_{j \in S} x_j p_j$.

To show this we use the PSD constraint on $X^{(i)}$. Let v be the $(|S|+1)$ dimensional vector indexed by $\emptyset$ and $\{ij\}_{j\in S}$ whose entries are defined by $v_\emptyset = -\mu$ and $v_{ij} = p_{ij} = p_j$ for $j \in S$. Let also $\overline{X}^{(i)}$ be the principal submatrix of $X^{(i)}$ containing those rows and columns indexed by $\emptyset$ and $\{ij\}_{j\in S}$. Then

$$\begin{aligned} v^T \overline{X}^{(i)} v &= X^{(i)}_{\emptyset,\emptyset} v_\emptyset^2 + 2\sum_{j\in S} X^{(i)}_{j,\emptyset} v_j v_\emptyset + \sum_{j,j'\in S} X^{(i)} X^{(i)}_{\{j\},\{j'\}} v_j v_{j'} \\ &= x_\emptyset \mu^2 - \sum_j 2 x_{\{j\}} p_j \mu + \sum_{j,j'\in S} x_{\{j\}\cup\{j'\}} p_j p_{j'} \\ &= \mu^2 - 2\mu^2 + \sum_{j,j'\in S} x_{\{j\}\cup\{j'\}} p_j p_{j'} = \sum_{j,j'\in S} p_{ij} p_{ij'} x_{\{ij\}\cup\{ij'\}} - \mu^2, \end{aligned}$$

which is greater than 0 because of the constraint $X^{(i)} \succeq 0$ in our relaxation (and hence the submatrix $\overline{X}^{(i)}$ is also positive semidefinite). This shows that $(II) \geq \frac{1}{2}\left(\sum_{j\in S} x_j p_j^2 + \left(\sum_{j\in S} x_j p_j\right)^2\right)$ and completes the proof of the lemma. □

Let $\mathrm{LB}(S)$ denote $\sum_{j\notin S} x_j p_j^2 + \frac{1}{2}\left(\sum_{j\in S} x_j p_j^2 + \left(\sum_{j\in S} x_j p_j\right)^2\right)$. By setting $S = \emptyset$, G and J, the lemma directly implies the following lower bounds:

Corollary 5.5. *We have the following lower bounds on* $\sum_{j=1}^{n'} p_j(p_1 x_{\{j\}\cup\{1\}} + \cdots + p_j x_{\{j\}\cup\{j\}})$:

$$\begin{aligned} \mathrm{LB}(\emptyset) &= Q, \\ \mathrm{LB}(J) &= 1/2(Q + L^2), \\ \mathrm{LB}(G) &= 1/2(\overline{Q} + Q + (L - \overline{L})^2). \end{aligned}$$

5.2.4 Proof of Inequality (5): bounding the approximation guarantee

We use Lemma 5.3 and Corollary 5.5 to prove (5). Let $\epsilon = 1/100$. We divide the proof into two cases. Intuitively, the first case is when we have "few" ungrouped jobs and then we get an improvement from the ζ in Theorem 1.2. In the other case, when we have "many" ungrouped jobs, note that jobs of different job classes have (informally) very different processing times and thus does not affect each other. This together with that the total fractional mass of ungrouped jobs in each class is less than $1/10$ actually gives that a simple randomized rounding does better than the factor $3/2$. The formal proof of the two cases are as follows:

Case $\overline{L} \leq (1-\sqrt{\epsilon})L$: We will upper bound the LHS of (5))

$$\left(1 - \epsilon\frac{\zeta}{100}\right)\mathrm{LB}(J) + \left(\frac{1}{2} - \frac{\zeta}{100} + \epsilon\frac{\zeta}{200}\right)\mathrm{LB}(\emptyset) + \frac{\zeta}{100}\mathrm{LB}(G). \tag{6}$$

By Corollary 5.5 this is at most

$$\left(1-\epsilon\frac{\zeta}{100}\right)+\left(\frac{1}{2}-\frac{\zeta}{100}+\epsilon\frac{\zeta}{200}\right)+\frac{\zeta}{100}=\left(\frac{3}{2}-\epsilon\frac{\zeta}{200}\right)$$

times the RHS of (5).

By the definition of $\mathrm{LB}(J), \mathrm{LB}(\emptyset)$ and $\mathrm{LB}(G)$, (6) can be written as

$$\begin{aligned}
&\left(1-\epsilon\frac{\zeta}{100}\right)\left(\frac{Q}{2}+\frac{L^2}{2}\right)+\left(\frac{1}{2}-\frac{\zeta}{100}+\epsilon\frac{\zeta}{200}\right)Q+\frac{\zeta}{100}\left(\frac{\overline{Q}+Q+(L-\overline{L})^2}{2}\right)\\
&=\left(1-\frac{\zeta}{200}\right)Q+\frac{\zeta}{200}\overline{Q}+\left(1-\epsilon\frac{\zeta}{100}\right)\frac{L^2}{2}+\frac{\zeta}{100}\frac{(L-\overline{L})^2}{2}\\
&\geq\left(1-\frac{\zeta}{200}\right)Q+\frac{\zeta}{200}\overline{Q}+\left(1-\epsilon\frac{\zeta}{100}\right)\frac{L^2}{2}+\epsilon\frac{\zeta}{100}\frac{L^2}{2}\qquad\text{(since }\overline{L}\leq(1-\sqrt{\epsilon})L\text{)}\\
&=\left(1-\frac{\zeta}{200}\right)Q+\frac{\zeta}{200}\overline{Q}+\frac{L^2}{2},
\end{aligned}$$

which is the upper bound on the LHS of (5) from Lemma 5.3 and thus completes this case.

Case $\overline{L}>(1-\sqrt{\epsilon})L$: Let $\overline{\mu}=\overline{L}/(\sum_{j\in\overline{G}}x_j)$ denote the expected job size in $\overline{G}$, i.e., of the ungrouped jobs, and let k denote the class of $\overline{\mu}$. Let $N\subseteq\overline{G}$ denote jobs in $\overline{G}$ in classes $k-1$ and higher. Also let $X(N)=\sum_{j\in N}x_j$.

We claim that $x(N)\leq 1/2$. Indeed, by Markov's inequality, the total mass of jobs in $\overline{G}$ in classes $k+2$ or higher is at most $1/10$. Moreover, as the mass of each class in $\overline{G}$ is at most $1/10$, we get $x(N)\leq 3/10+1/10\leq 1/2$.

Let us define $L(N)=\sum_{j\in N}x_jp_j$ and $Q(N)=\sum_{j\in N}x_jp_j^2$. By Cauchy-Schwarz, we have $(\sum_{j\in N}x_jp_j^2)(\sum_{j\in N}x_j)\geq(\sum_{j\in N}x_jp_j)^2$ and hence

$$Q(N)\geq\frac{L(N)^2}{x(N)}\geq 2L(N)^2.$$

Next we show that the total expected size of jobs in $\overline{G}\setminus N$ is negligible compared to $L(N)$. Indeed a job of class $k-h$ has processing time at most $\overline{\mu}/10^{h-1}$ and the total mass of jobs of class $k-h$ in $\overline{G}$ is at most $\sum_{j\in\overline{G}}x_j$ since this is the total mass of all jobs in $\overline{G}$. This gives us the rough upper bound

$$\overline{L}-L(N)\leq\sum_{h=2}^{\infty}\frac{1}{10^{h-1}}\cdot\overline{\mu}\left(\sum_{j\in\overline{G}}x_j\right)=\frac{\overline{L}}{9}.$$

The above gives us that

$$\begin{aligned}
L^2&\leq\frac{\overline{L}^2}{(1-\sqrt{\epsilon})^2}\leq\left(\frac{9}{8(1-\sqrt{\epsilon})}\right)^2L(N)^2\leq\left(\frac{9}{8(1-\sqrt{\epsilon})}\right)^2\frac{Q(N)}{2}\\
&=\left(\frac{10}{8}\right)^2\frac{Q(N)}{2}=\frac{25}{32}Q(N)\leq\frac{25}{32}Q.
\end{aligned}$$

By Lemma 5.3, we have that the LHS of (5) is at most [4]

$$Q + \frac{L^2}{2} \leq \left(1 + \frac{25}{64}\right) Q = \left(1 + \frac{25}{64}\right) \mathrm{LB}(\emptyset) < \left(\frac{3}{2} - c\right) \mathrm{LB}(\emptyset),$$

which completes this case and the proof of (5) (and thus Theorem 5.1).

Acknowledgement

This work was done in part while the first author was visiting the Simons Institute for the Theory of Computing. We thank Nick Harvey and Bruce Shepherd for organizing the Bellairs Workshop on Combinatorial Optimization 2015, which was the starting point for this work. Aravind Srinivasan thanks Amit Chavan for his help with LaTeXpackages.

References

[1] Foto N. Afrati, Evripidis Bampis, Chandra Chekuri, David R. Karger, Claire Kenyon, Sanjeev Khanna, Ioannis Milis, Maurice Queyranne, Martin Skutella, Clifford Stein, and Maxim Sviridenko. Approximation schemes for minimizing average weighted completion time with release dates. In *Foundations of Computer Science, FOCS*, pages 32–44, 1999.

[2] Alexander A. Ageev and Maxim Sviridenko. Approximation algorithms for maximum coverage and max cut with given sizes of parts. In *Integer Programming and Combinatorial Optimization IPCO*, pages 17–30, 1999.

[3] Sanjeev Arora, Alan M. Frieze, and Haim Kaplan. A new rounding procedure for the assignment problem with applications to dense graph arrangement problems. *Math. Program.*, 92(1):1–36, 2002.

[4] Arash Asadpour, Uriel Feige, and Amin Saberi. Santa claus meets hypergraph matchings. *ACM Transactions on Algorithms*, 8(3):24, 2012.

[5] Arash Asadpour and Amin Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. In *Symposium on Theory of Computing, STOC*, pages 114–121, 2007.

[6] Yossi Azar and Amir Epstein. Convex programming for scheduling unrelated parallel machines. In *Symposium on Theory of Computing*, pages 331–337, 2005.

[7] Nikhil Bansal and Maxim Sviridenko. The santa claus problem. In *Symposium on Theory of Computing, STOC*, pages 31–40, 2006.

[8] Deeparnab Chakrabarty, Julia Chuzhoy, and Sanjeev Khanna. On allocating goods to maximize fairness. In *Foundations of Computer Science, FOCS*, pages 107–116, 2009.

[9] Chandra Chekuri and Sanjeev Khanna. A PTAS for minimizing weighted completion time on uniformly related machines. In *ICALP*, pages 848–861, 2001.

[10] Chandra Chekuri and Sanjeev Khanna. Approximation algorithms for minimizing averageweighted completion time. In *Handbook of Scheduling - Algorithms, Models, and Performance Analysis.* 2004.

[4] This bound is loose and can also be obtained using a independent randomized rounding instead of our negative correlation rounding. The importance of the new rounding appears in the other case.

[11] Chandra Chekuri, Rajeev Motwani, B. Natarajan, and Clifford Stein. Approximation techniques for average completion time scheduling. *SIAM J. Comput.*, 31(1):146–166, 2001.

[12] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Multi-budgeted matchings and matroid intersection via dependent rounding. In *Symposium on Discrete Algorithms, SODA*, pages 1080–1097, 2011.

[13] Uriel Feige. On allocations that maximize fairness. In *Symposium on Discrete Algorithms, SODA*, pages 287–293, 2008.

[14] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. *J. ACM*, 53(3):324–360, 2006.

[15] Michel X. Goemans, Maurice Queyranne, Andreas S. Schulz, Martin Skutella, and Yaoguang Wang. Single machine scheduling with release dates. *SIAM J. Discrete Math.*, 15(2):165–192, 2002.

[16] Leslie A. Hall, Andreas S. Schulz, David B. Shmoys, and Joel Wein. Scheduling to minimize average completion time: Off-line and on-line approximation algorithms. *Mathematics of Operations Research*, 22(3):513–544, 1997.

[17] Leslie A. Hall, David B. Shmoys, and Joel Wein. Scheduling to minimize average completion time: Off-line and on-line algorithms. In *Symposium on Discrete Algorithms, SODA*, pages 142–151, 1996.

[18] Han Hoogeveen, Petra Schuurman, and Gerhard J. Woeginger. Non-approximability results for scheduling problems with minsum criteria. In *Integer Programming and Combinatorial Optimization, IPCO*, pages 353–366, 1998.

[19] Jeff Kahn and P. Mark Kayll. On the stochastic independence properties of hard-core distributions. *Combinatorica*, 17(3):369–391, 1997.

[20] V. S. Anil Kumar, Madhav V. Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. Minimum weighted completion time. In *Encyclopedia of Algorithms*. 2008.

[21] V. S. Anil Kumar, Madhav V. Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. A unified approach to scheduling on unrelated parallel machines. *J. ACM*, 56(5), 2009.

[22] Jan Karel Lenstra, David B. Shmoys, and Éva Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Math. Program.*, 46:259–271, 1990.

[23] Konstantin Makarychev and Maxim Sviridenko. Solving optimization problems with diseconomies of scale via decoupling. In *Foundations of Computer Science, FOCS*, pages 571–580, 2014.

[24] Cynthia A. Phillips, Clifford Stein, and Joel Wein. Task scheduling in networks. *SIAM J. Discrete Math.*, 10(4):573–598, 1997.

[25] Cynthia A. Phillips, Clifford Stein, and Joel Wein. Minimizing average completion time in the presence of release dates. *Math. Program.*, 82:199–223, 1998.

[26] Andreas S. Schulz and Martin Skutella. Scheduling unrelated machines by randomized rounding. *SIAM J. Discrete Math.*, 15(4):450–469, 2002.

[27] Petra Schuurman and Gerhard Woeginger. Polynomial time approximation algorithms for machine scheduling: Ten open problems. *Journal of Scheduling*, 2(5):203–213, 1999.

[28] Jay Sethuraman and Mark S. Squillante. Optimal scheduling of multiclass parallel machines. In *ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 963–964, 1999.

[29] Martin Skutella. Personal communication. Oct 2015.

[30] Martin Skutella. Convex quadratic and semidefinite programming relaxations in scheduling. *J. ACM*, 48(2):206–242, 2001.

[31] Martin Skutella and Gerhard J. Woeginger. A PTAS for minimizing the weighted sum of job completion times on parallel machines. In *Symposium on Theory of Computing, STOC*, pages 400–407, 1999.

[32] W. E. Smith. Various optimizers for single-stage production. *Naval Research Logistics*, 3:5966, 1956.

[33] Ola Svensson. Santa claus schedules jobs on unrelated machines. *SIAM J. Comput.*, 41(5):1318–1341, 2012.

[34] Maxim Sviridenko and Andreas Wiese. Approximating the configuration-lp for minimizing weighted sum of completion times on unrelated machines. In *Integer Programming and Combinatorial Optimization, IPCO*, pages 387–398, 2013.